

AMENDMENTS TO THE CLAIMS

1. (previously presented) A method for replicating data from a storage device, comprising:
  - identifying on a storage device at least one data block comprising file data;
  - performing at least one read operation with respect to the at least one data block;
  - recording one or more I/O accesses performed with respect to the storage device in association with the at least one read operation;
  - identifying, based on the recorded I/O access information, one or more data blocks on the storage device that contain valid data; and
  - replicating the data blocks that contain valid data.
2. (previously presented) The method according to claim 1, wherein the at least one read operation includes reading metadata associated with at least one file on the storage device.
3. (previously presented) The method according to claim 2, wherein reading the metadata includes reading one or more of the following: a name of a file, access permissions to a file, a date of creation of a file, or dates of modification of a file.
4. (original) The method according to claim 1, further comprising cleaning a cache on a computer associated with the storage device before performing any read operations.

5. (previously presented) A method for replicating data from a storage device associated with a computer, comprising:

cleaning a cache on a computer associated with a storage device;  
identifying on the storage device at least one data block comprising file data;  
performing at least one read operation with respect to the at least one data block;

causing the storage device to record one or more I/O accesses performed with respect to the storage device in association with the at least one read operation;

identifying, based on the I/O access information recorded by the storage device, one or more data blocks on the storage device that contain valid data; and

replicating the data blocks that contain valid data.

6. (cancelled)

7. (cancelled)

8. (cancelled)

9. (cancelled)

10. (previously presented) A system to identify data blocks on a storage device that contain valid data, comprising:

a storage device configured to:

store data in one or more data blocks;  
at least one processor configured to:  
identify at least one data block comprising file data;  
perform at least one read operation with respect to the at least one data  
block;  
record one or more I/O accesses performed with respect to the storage  
device in association with the at least one read operation;  
identify one or more data blocks on the storage device that contain  
valid data based, at least in part, on the recorded I/O access information; and  
replicate the data blocks that contain valid data.

11. (previously presented) The system according to claim 10, wherein the at least one read operation includes reading metadata associated with at least one file on the storage device.

12. (previously presented) The system according to claim 11, wherein reading the metadata includes reading one or more of the following: a name of a file, access permissions to a file, a date of creation of a file, or dates of modification of a file.

13. (original) The system according to claim 10, further comprising a computer associated with the storage device.

14. (previously presented) The system according to claim 13, wherein the at least one processor resides on the computer.

15. (previously presented) The system according to claim 13, wherein the at least one processor is further configured to:

clean a cache on the computer before performing any I/O accesses.

16. (previously presented) The system according to claim 13, wherein the at least one processor is further configured to:

manage the storage operations of the computer.

17. (previously presented) The system according to claim 10, wherein the at least one processor comprises a filter driver.

18. (currently amended) An apparatus to ~~identify~~ replicate data blocks on a storage device that contain valid data, comprising:

a storage device configured to:

store data in one or more data blocks;

wherein the storage system comprises a file system that identifies files stored on the storage device and storage location information for the respective files;

a first processor configured to:

record I/O accesses performed with respect to the storage device in association with read operations; and

a second processor configured to:

~~identify at least one data block comprising file data;~~

~~perform at least one read operation with respect to the at least one data block;~~

perform read operations with respect to all files identified in the file system; and

instruct the first processor to record one or more I/O accesses performed with respect to the storage device in association with the ~~at least one~~ read operations;

wherein the first processor is further configured to:

identify one or more data blocks on the storage device that contain valid data based, at least in part, on the I/O access information recorded by the first processor; and

replicate the data blocks that contain valid data.

19. (previously presented) The apparatus according to claim 18, wherein the second processor is further configured to:

clean a cache on a computer associated with the storage device before performing any I/O accesses.

20. (currently amended) The apparatus according to claim 18, wherein ~~the~~ at least one of the read operations includes reading metadata associated with at least one file on the storage device.

21. (previously presented) The apparatus according to claim 20, wherein reading the metadata includes reading one or more of the following: a name of a file, access permissions to a file, a date of creation of a file, or dates of modification of a file.

22. (previously presented) The apparatus according to claim 18, wherein the second processor comprises a software program.

23. (previously presented) The apparatus according to claim 18, wherein the second processor comprises a filter driver.

24. (previously presented) The apparatus according to claim 18, wherein the second processor is part of a storage management system.

25. (currently amended) The method of claim 1, wherein ~~the~~ a file system associated with the storage device is structured on a file-level, the method comprising:  
identifying on the storage device at least one data block comprising file data, based at least in part, on information in the file system.

26. (previously presented) The method of claim 5, wherein the at least one read operation includes reading metadata associated with one or more files on the storage device.

27. (previously presented) The method according to claim 26, wherein reading the metadata includes reading one or more of the following: a name of a file, access permissions to a file, a date of creation of a file, or dates of modification of a file.

28. (previously presented) The method according to claim 1, further comprising:  
generating a list of the one or more data blocks that contain valid data; and  
storing the list and the replicated data blocks in a memory.

29. (currently amended) A method to identify data blocks on a storage device that contain valid data, comprising:

identifying on a storage device at least one data block comprising file data;  
performing at least one read operation with respect to at least one data block;  
recording one or more I/O accesses performed with respect to the storage device in association with the at least one a read operation; and  
generating a single list of all data blocks on the storage device that contain valid data based, at least in part, on the recorded I/O access information.

30. (currently amended) The method of claim 29, wherein: the storage device comprises:

~~the file system is associated with~~ a virtual storage device used to manage storage of data on the storage device; and

a file system associated with the virtual storage device;

the method further comprising:

identifying at least one data block comprising file data based at least in part on the file system.

31. (currently amended) The method of claim 29, further comprising:  
storing the single list in a memory.

32. (previously presented) A system to identify data blocks on a storage device that contain valid data, comprising:

a storage device configured to:

store data in one or more data blocks;

at least one processor configured to:

identify at least one data block comprising file data;

perform at least one read operation with respect to the at least one data block;

record one or more I/O accesses performed with respect to the storage device in association with a the at least one read operation; and

generate a list of data blocks on the storage device that contain valid data based at least in part on the recorded I/O access information.

33. (currently amended) The system of claim 32, wherein: the storage device comprises:

~~the file system is associated with~~ a virtual storage device used to manage storage of data on the storage device; and



a file system associated with the virtual storage device;

the at least one processor being configured to:

identify at least one data block comprising file data based at least in part on the file system.

34. (previously presented) The system of claim 32, wherein the at least one processor is further configured to:

store the list in a memory.

35. (previously presented) The system of claim 10, wherein the at least one processor comprises a first processor and a second processor, wherein:

the first processor is configured to:

identify at least one data block comprising file data;

perform at least one read operation with respect to the at least one data block;

identify one or more data blocks on the storage device that contain valid data based, at least in part, on the recorded I/O access information; and

replicate the data blocks that contain valid data; and

the second processor is configured to:

record one or more I/O accesses performed with respect to the storage device in association with the at least one read operation.

36. (previously presented) The system of claim 10, wherein the at least one processor comprises at least one computer.

37. (previously presented) The system of claim 36, wherein the at least one processor comprises:

a first software program operating on the at least one computer; and

a second software program operating on the at least one computer.

38. (previously presented) A method for replicating data from a storage device, comprising:

receiving a message to replicate data stored on a storage device;

in response to the message, identifying on the storage device at least one data block comprising file data;

performing at least one read operation with respect to the at least one data block;

storing in a file information identifying the at least one data block; and

replicating one or more data blocks stored on the storage device, based, at least in part, on the information in the file.

39. (previously presented) The method of claim 38, further comprising:

recording, in the file, one or more I/O accesses performed by the storage device.

40. (previously presented) The method of claim 39, wherein the message comprises a request to record I/O accesses performed by the storage device.

41. (previously presented) The method of claim 40, further comprising:

receiving a second message to stop recording I/O accesses performed by the storage device.

42. (previously presented) A system for replicating data from a storage device, comprising:

a storage device configured to:

store data;

a first processor configured to:

transmit a first message to replicate data stored on the storage device; and

identify one or more data blocks comprising file data;

read each identified data block, until all data blocks containing valid data have been identified and read;

a second processor configured to:

receive the message;

record one or more I/O accesses performed by the storage device while the identified data blocks are being read, in response to the message; and

replicate at least one data block stored on the storage device, based on the recorded I/O accesses.

43. (previously presented) The method of claim 1, comprising:

identifying on the storage device at least one data block comprising file data and referenced in a file system associated with the storage device.

44. (previously presented) The method of claim 43, further comprising:  
examining the file system; and  
identifying one or more files stored on the storage device based on information  
in the file system.

45. (new) A method to replicate data stored in a storage system,  
comprising:  
examining a file system associated with the storage system, wherein the file  
system specifies one or more files and identifies one or more storage locations associated with  
at least one of the one or more files;  
reading each file specified in the file system; and  
recording in a list an identifier of at least one storage location accessed in  
association with each read operation; and  
after at least one read operation is performed with respect to each file identified  
in the file system:  
replicating data stored in each storage location identified in the list.